

CLAIMS

1. A method of automatically configuring a microprocessor architecture, comprising the step of using executable code for another type of microprocessor.
2. The method according to claim 1 whereby the configured processor is used as a coprocessor in a system.
3. The method according to claim 2 whereby the host processor in the system is of the type of the executable used to configure the coprocessor.
4. The method according to claim 3 whereby the host processor in a system runs the executable used to configure the coprocessor.
5. The method according to claim 4 whereby a number of individual software functions in the executable code are marked for translation and execution on the coprocessor.
6. The method according to claim 5 whereby the original executable image is automatically modified so that function calls to those translated functions cause an equivalent function to be executed on the coprocessor.
7. The method according to claim 6 whereby the coprocessor initiation involves the transfer of register state from the host processor to the coprocessor.
8. The method according to claim 7 whereby the completion of a function on the coprocessor causes the transfer of register state from the coprocessor to the host processor.
9. The method according to claim 7 whereby the completion of a function on the coprocessor causes the transfer of memory state from the coprocessor to the host processor.
10. The method according to claim 1 whereby the architecture generated is designed to execute parts of the executable with higher performance than can be achieved with the host processor.

11. The method according to claim 10 whereby the improved performance is obtained by the execution of more operations in parallel than is achieved with the host processor.
12. The method according to claim 1 whereby the architecture generated is designed to execute parts of the executable with lower power consumption than can be achieved with the host processor.
13. The method according to claim 1 whereby the executable code is translated into the instruction set of the configured processor.
14. The method according to claim 13 whereby each instruction in the executable image is translated into one or more basic operations.
15. The method according to claim 14 whereby each of these operations may be performed by particular execution unit that is present in the configured processor.
16. The method according to claim 15 whereby the register file is present as an execution unit in the architecture and explicit operations to read and write the register file is generated as part of the translation.
17. The method according to claim 16 whereby static register analysis may be used to eliminate unnecessary writes of registers.
18. The method according to claim 17 whereby code is subdivided into atomically executed blocks.
19. The method according to claim 18 whereby each atomically executed block reproduces the operations of the corresponding host code.
20. The method according to claim 19 whereby the state of live registers at the end of the atomic block execution is identical to that obtained from execution on the host processor.
21. The method according to claim 19 whereby the state of the memory at the end of the atomic block execution is identical to that obtained from execution on the host processor.

22. The method according to claim 1 whereby breakpoints may be set on the configured architecture.
23. The method according to claim 2 whereby the breakpoints may be specified using the addresses of instructions in the original executable.
24. The method according to claim 3 whereby the nearest preceding instruction for which state can be synchronised on the configured processor is determined when the breakpoint is set.
25. The method according to claim 24 whereby the configured processor contains a mechanism to determine the equivalent target instruction address for a host instruction address.
26. The method according to claim 25 whereby the configured processor contains hardware to cause a breakpoint halt on the required address that prevents any side effects caused by sequentially later instructions.
27. The method according to claim 26 whereby upon detection of a breakpoint execution can be continued on the host processor from the synchronised address until the point of the actual breakpoint.
28. The method according to claim 23 whereby the breakpoint address is determined by decoding the data stream on the debug interface to the host processor.
29. The method according to claim 1 whereby certain host processor instruction addresses may be converted to target processor addresses while the system is running.
30. The method according to claim 29 whereby a hashing table is maintained in memory to perform a mapping of certain host processor instruction addresses to target addresses.
31. The method according to claim 30 whereby the mapping information may be interleaved with the machine code for the target processor.

32. The method according to claim 31 whereby the state of certain bits within each word of a table are used to indicate whether the information represents an address mapping entry or target machine code.
33. The method according to claim 1 whereby function calls in the executable code may be replaced with uses of particular hardware blocks in the configured processor.
34. The method according to claim 33 whereby the input parameters to the software function correspond to the operands supplied to the corresponding hardware unit.
35. The method according to claim 34 whereby the reference parameters and return result from a software function correspond to the results generated by the corresponding hardware unit.
36. The method according to claim 33 whereby the original software implementation may be used as a behavioural model for hardware for the purposes of simulation.
37. The method according to claim 1 whereby an instruction set translator converts instructions from the executable image into behaviourally equivalent operations that are mapped to the target processor.
38. The method according to claim 4 whereby the coprocessor contains one or more cache memories.
39. The method according to claim 38 whereby the coprocessor and host processor communicate via a system bus or a generic coprocessor interface on the host processor.
40. The method according to claim 39 whereby the host processor services memory access requests from the coprocessor while the coprocessor is operational.
41. The method according to claim 40 whereby the coprocessor is able to flush its caches of all modified data when the end of function execution is reached.
42. The method according to claim 39 whereby a copy of some virtual to physical page mappings are maintained by the coprocessor.

43. A microprocessor that has been automatically configured using the method as defined in any preceding claim 1 – 42.